Mon. Not. R. Astron. Soc. **000**, 1–8 (2012) Printed 3 March 2014 (MN LATEX style file v2.2)

Neutron star glitches have a substantial minimum size

C.M. Espinoza,[1,2⋆] D. Antonopoulou,[3] B.W. Stappers,[1] A. Watts[3] and A.G. Lyne[1]
[1] *Jodrell Bank Centre for Astrophysics, School of Physics and Astronomy, The University of Manchester, Manchester M13 9PL, UK.*
[2] *Instituto de Astrofísica, Facultad de Física, Pontificia Universidad Católica de Chile, Casilla 306, Santiago 22, Chile.*
[3] *Astronomical Institute Anton Pannekoek, University of Amsterdam, Postbus 94249, 1090GE Amsterdam, The Netherlands.*

3 March 2014

arXiv:1402.7219v1 [astro-ph.HE] 28 Feb 2014

ABSTRACT
Glitches are sudden spin-up events that punctuate the steady spin down of pulsars and are thought to be due to the presence of a superfluid component within neutron stars. The precise glitch mechanism and its trigger, however, remain unknown. The size of glitches is a key diagnostic for models of the underlying physics. While the largest glitches have long been taken into account by theoretical models, it has always been assumed that the minimum size lay below the detectability limit of the measurements. In this paper we define general glitch detectability limits and use them on 29 years of daily observations of the Crab pulsar, carried out at Jodrell Bank Observatory. We find that all glitches lie well above the detectability limits and by using an automated method to search for small events we are able to uncover the full glitch size distribution, with no biases. Contrary to the prediction of most models, the distribution presents a rapid decrease of the number of glitches below $\sim 0.05\,\mu$Hz. This substantial minimum size indicates that a glitch must involve the motion of at least several billion superfluid vortices and provides an extra observable which can greatly help the identification of the trigger mechanism. Our study also shows that glitches are clearly separated from all the other rotation irregularities. This supports the idea that the origin of glitches is different to that of timing noise, which comprises the unmodelled random fluctuations in the rotation rates of pulsars.

Key words: pulsars: general – pulsars: individual: PSR B0531+21 – stars: neutron

1 INTRODUCTION

Neutron stars are the highly-magnetised and rapidly-rotating remnants of the collapse of the cores of once more-massive stars. Having masses of approximately $1.4\,M_\odot$ and radii of about 12 km, the high densities of neutron stars indicate a structure of a crystalline-like crust and a superfluid interior (Baym et al. 1969; Haensel et al. 2007). Their large and steady moments of inertia mean that they have extremely stable rotational frequencies, which slowly decrease as energy is lost through electromagnetic radiation and acceleration of particles in their magnetospheres. However, this regular spin-down is occasionally interrupted by sudden spin-up events, known as glitches (Radhakrishnan & Manchester 1969; Espinoza et al. 2011b).

The exact mechanism responsible for glitches is not fully understood but it is thought to involve a sudden transfer of angular momentum from a more rapidly rotating superfluid component to the rest of the star (Anderson & Itoh 1975). This component resides in regions of the interior where neutron vortices, which carry the angular momentum of the superfluid, are impeded in moving by pinning on crustal nuclei or on superconducting vortices in the core (or on both). Since a superfluid in such conditions cannot slow down by outwards motion and expulsion of vortices, the superfluid component will retain a higher rotational frequency as the rest of the star slows down. A glitch occurs when vortices are suddenly unpinned and free to move outwards, allowing for a rapid exchange of angular momentum and the observed spin-up of the crust.

Catastrophic unpinning of vortices is expected once the velocity lag between the two components exceeds a maximum threshold, above which the pinning force can no longer sustain the hydrodynamic lift force exerted on the pinned vortices by the ambient superfluid. It has also been shown (Glampedakis & Andersson 2009; Andersson et al. 2013) that, beyond some critical lag, a two-stream instability might develop and trigger the unpinning. If in such events the lag is completely relaxed (or partially relaxed by a fixed amount) then the interglitch time interval corresponds to the time it takes for the system to reach the critical threshold again, driven by the nearly constant external torque.

⋆ E-mail: cespinoz@astro.puc.cl

© 2012 RAS

Models relying on such a build-up and depletion of the superfluid angular momentum reservoir have been successfully used to explain the regular, similar glitches of some young pulsars (Alpar et al. 1993; Pizzochero 2011; Haskell et al. 2012). However this simple picture cannot account for the wide range of glitch sizes and waiting times between glitches seen in most pulsars.

Glitch sizes in rotational frequency can range over four orders of magnitude in individual pulsars and appear to follow a power-law distribution (Melatos et al. 2008). This favours scale-invariant models of the dynamics of individual vortices in the presence of a pinning potential, such as the vortex avalanche model (Warszawski & Melatos 2008) and the coherent noise model (Melatos & Warszawski 2009). Alternative models involve non-superfluid mechanisms that can act as unpinning triggers before the critical lag is reached, such as crustquakes (Ruderman 1969; Baym & Pines 1971) or heating episodes (Link & Epstein 1996). The crustquake-induced glitch model has been particularly favoured for the Crab pulsar as it may explain the persistent changes in slow-down rate observed after some of its glitches (Gullahorn et al. 1977; Alpar et al. 1994) and could possibly lead to a power-law distribution of event sizes, similar to earthquakes.

A second type of irregularity is often seen in the rotational behaviour of pulsars, namely timing noise. Thought to be partially caused by torque variations driven by two or more magnetospheric states (Lyne et al. 2010), it manifests as a continuous and erratic wandering of the rotation rate around the predictions of a simple slow-down model. While glitches are rapid and sporadic events in rotation rate, timing noise appears as a slow and continuous process.

Owing to observational limitations such as infrequent and irregular sampling and the presence of timing noise, the detection of glitches is an uncertain process. Moreover, the signature of timing noise in the data can be confused with glitches, so that the lower end of a glitch-size distribution is possibly contaminated by spurious detections. Knowledge of this distribution is essential for any glitch theory. The largest glitches are easily detected and can be used to constrain the minimum superfluid moment of inertia that can act as an angular momentum reservoir (Andersson et al. 2012; Chamel 2013). The biases involved and the question of whether there is a minimum glitch size have not been addressed; so far, the smallest possible glitch has been assumed to lie below our detection limits.

In this paper we study the glitch detection capabilities of the current detection methods and define limits depending upon the intrinsic pulsar rotational stability, observing cadence and sensitivity. We apply these definitions to an extensive set of observations of the Crab pulsar and, by using an automated glitch detector, uncover the full glitch size distribution and show that there is a minimum glitch size.

2 LIMITS ON GLITCH DETECTION

To assess the level of completeness of the existing glitch samples, we quantify simple observational limits on glitch detection, applicable for a given pulsar and observing setup. The first step towards this is establishing a working definition of what constitutes a glitch. Traditionally, glitches are identified by visual inspection of the pulsar's timing residuals, which are defined as the phase differences between measurements and the predictions of a model for the rotation. To put this on a more formal footing, we define a glitch as an event characterised by a sudden, discrete positive change in rotational frequency ($\Delta\nu$) and a discrete negative or null change in frequency spin-down rate ($\Delta\dot{\nu}$). These two sudden changes together make glitches distinguishable from timing noise (Espinoza et al. 2011b; Lyne et al. 2010).

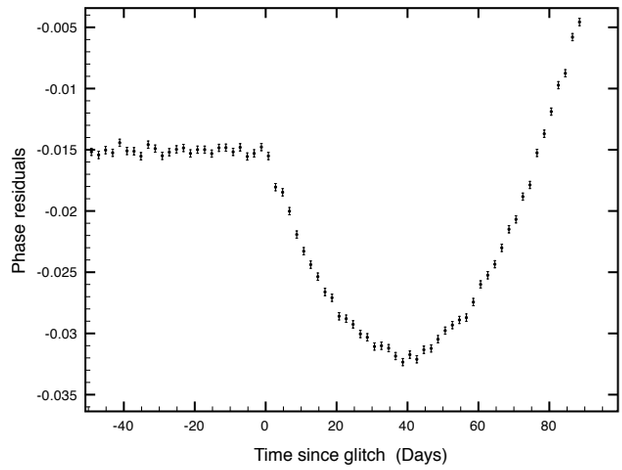

Figure 1. Example of a glitch signature in the timing residuals. The residuals are with respect to a model which describes well the rotation before the glitch ($t < 0$ on the plot). This is a simulated glitch in the Crab pulsar's rotation, with $\Delta\nu = 0.01\,\mu$Hz and $\Delta\dot{\nu} = -3.0 \times 10^{-15}\,\mathrm{Hz\,s^{-1}}$.

The timing residuals will be flat if the model describes the rotation of the pulsar well. For such a model, the timing residuals after a glitch at $t = t_g$ will follow a quadratic signature given by

$$\phi_g = -\Delta\nu(t - t_g) - \Delta\dot{\nu}\frac{(t - t_g)^2}{2} \quad ; \quad (t > t_g)\,. \qquad (1)$$

The frequency change $\Delta\nu > 0$ produces a linear drift of the post-glitch residuals towards negative values, with the slope being the magnitude of the frequency step. The effect of a change $\Delta\dot{\nu} < 0$ is a parabolic signature which lifts the residuals towards positive values. Therefore a glitch with a large, negative change in spin-down rate will produce positive residuals rising quadratically soon after the glitch (Fig. 1).

Based upon these facts, we can define simple limits that describe our ability to detect glitches in the timing residuals. If the observing cadence is not very frequent, it is possible that no observations occur before the post-glitch residuals rise above the extrapolation of the line defined by the pre-glitch ones (Fig. 1). This effect will primarily mask glitches with small $\Delta\nu$ and large $|\Delta\dot{\nu}|$. Requiring at least one observation before the rise of the residuals defines a minimum $\Delta\nu$ that can be detected (Eq. 2). If $\Delta\nu$ became smaller, the dip would become shallower and in the case that it is undetectably small, the event is unlikely to be recognised as a glitch and might appear as timing noise. To ensure detection, the maximum negative departure of the residuals ought to be larger than both the root mean square (RMS) of the

timing residuals prior to the glitch and the typical error of the TOAs. Therefore, a detectable glitch is a rapid event in which the effects of $\Delta\nu$ are recognisable over the effects of $\Delta\dot{\nu}$ and the limiting detectable value of $\Delta\nu$ depends on the observation cadence (one observation every ΔT days) and the largest of either the sensitivity of the observations or the typical dispersion of the timing residuals in rotational phase, σ_ϕ, as

$$\Delta\nu_{\rm lim} = \max \begin{cases} \Delta T|\Delta\dot{\nu}|/2 \\ \\ \sqrt{2\,\sigma_\phi|\Delta\dot{\nu}|} \end{cases} \qquad (\text{for} \quad \Delta\dot{\nu} < 0). \qquad (2)$$

For simplicity and because of our particular focus on small events, any exponential recovery of the frequency, often observed after glitches, is not considered here. Nonetheless, we constructed several detectability curves, with decaying components and timescales similar to those observed in the Crab pulsar, and verified that our conclusions are not altered if exponential recoveries are present.

These limits are consistent with the glitch samples of several pulsars, hence we believe they offer a realistic way to assess glitch detectability as it is commonly carried out.

3 THE CRAB PULSAR GLITCHES

The Crab pulsar (PSR B0531+21; PSR J0534+2200) is the central source of the Crab Nebula and a young neutron star widely studied since it was first observed in 1968. The rotation of the Crab pulsar has been monitored almost every day for the last 29 years with the 42-ft radio telescope operating at 610 MHz at the Jodrell Bank Observatory (JBO) in the UK (Lyne et al. 1988, 1993). This offers an ideal dataset to test the completeness of the glitch sample because of its rapid cadence, good sensitivity and low dispersion of the timing residuals.

3.1 Observations

The product of each observation was the time of arrival (TOA) of one pulse at the observatory, corrected to the solar system barycenter. The dataset comprises 8862 TOAs starting in January 1984. There is one TOA per day in general and two TOAs per day during some periods of time. In addition, towards the beginning of the dataset, there are some isolated cases in which groups of TOAs are separated by up to 5 days. Finally, there are also a few gaps with no observations, generally no larger than ~ 20 days, when the telescope or observing hardware were unavailable due to maintenance.

The TOAs generally have errors of less than 0.001 rotation, with more than 75% having uncertainties less than 0.0004 rotation. For groups of 20 TOAs, which cover 20 days on average, the timing residuals with respect to a simple slow-down model with two frequency derivatives typically give a dispersion similar to the TOA uncertainties (hence $\sigma_\phi \sim 0.0004$ rotations).

3.2 Detection limits and the sample of detected glitches

To study the glitch size distribution of the Crab pulsar we need a complete list of glitches for the time interval defined by the 42-ft dataset, and their main parameters $(\Delta\nu, \Delta\dot{\nu})$. As described above, we classify events as glitches based on the assumption that a glitch is a sudden, unresolved change in spin frequency, implying clearly defined features in the timing residuals.

We use the events included in the JBO online glitch catalogue[1], which correspond to all the events published by Espinoza et al. (2011b) plus one new glitch that occurred on MJD 55875.5 (Espinoza et al. 2011a). The event on MJD $\sim$ 50489, originally reported by Wong et al. (2001), was rejected because of its anomalous characteristics, already described by them. No other glitches have been reported for this time-span by other authors (e.g. Wong et al. 2001; Wang et al. 2012) and we confirmed this by visually inspecting the timing residuals for all our dataset. Our final list contains 20 glitches, with parameters covering the ranges $0.05 \leq (\Delta\nu/\mu{\rm Hz}) \leq 6.37$ and $45 \leq \left(|\Delta\dot{\nu}|/10^{-15}{\rm Hz\,s^{-1}}\right) \leq 2302$. Here we use the glitch sizes reported by Espinoza et al. (2011b,a).

We note the clear presence of four other glitches prior to the start of the 42-ft observations. However, the available data for that period is highly inhomogeneous and contains large gaps with no observations, making it difficult to define single detectability limits and complicating the use of the glitch detector (see below). Hence, in order to work with a set of glitches that we know is statistically complete, we have not included them in our sample.

Using $\Delta T = 1\,$day and $\sigma_\phi = 0.0004$ in Eq. 2 we find that all glitches (including the four early ones) show a clear separation from the detectability limits (Fig. 2). Thus, at least for intermediate and large glitches, we are uncovering the true $\Delta\nu$–$\Delta\dot{\nu}$ distribution, with no biases.

4 THE GLITCH DETECTOR

To confirm that we have identified every glitch in the data, especially small ones which may be missed by standard techniques, we developed an automated glitch detector to find and measure every timing signature that might be regarded as a glitch. The detector assumes that a glitch occurred after every observation and attempts to measure its size, producing an output of glitch candidates (GCs) whenever $\Delta\nu > 0$ and $\Delta\dot{\nu} \leq 0$ are detected.

4.1 Method

The detector's technique is based on the fact that timing residuals, in the presence of timing noise or glitches, quickly depart from the best fit model of previous data, resulting in deviations from a mean of zero as newer observations are included and the model is not updated. Below we describe the method step by step, optimised for the JBO dataset for the Crab pulsar, described above. Different parameters should be used for different datasets.

A fit for ν, $\dot{\nu}$ and $\ddot{\nu}$ is performed over a set of 20 TOAs using the timing software PSRTIME and TEMPO2 (Hobbs et al. 2006), following standard techniques. To test for a glitch occurring after the last TOA in a set, the timing residuals of

[1] `http://www.jb.man.ac.uk/pulsar/glitches.html`

the following 10 TOAs, relative to that model, are fitted with a quadratic function of the form of Eq. (1) and separately with just the linear term in that equation. The latter is to test the case $\Delta\dot{\nu} = 0$. The fit with the smallest reduced χ^2 is selected. When the quadratic fit is selected, an event is characterised as a GC only if the reduced χ^2 is less than 15, the quadratic part of the fit is negative ($\Delta\dot{\nu} < 0$) and if the minimum of the fitted curve is at least 2.5 times the dispersion of the timing residuals of the 20 TOAs below zero. This last condition ensures a positive $\Delta\nu$ and a solid detection of its magnitude. If the linear fit is selected, a new GC is created only if the reduced χ^2 is less than 15 and the slope of the fit is negative, indicating $\Delta\nu > 0$. In this case the GC has a null or undetectable $\Delta\dot{\nu}$. The conclusions of our analysis are not dependent on the choice of the maximum allowed reduced χ^2 threshold[2]. The chosen value of 15 is high enough to avoid missing signatures that one might regard as a glitch.

The next step is to move the analysis forward by one TOA to define a new set of 20 TOAs and test for a glitch occurring after this new TOA. By doing this over the whole dataset, the dataset is explored for glitches after every single observation (with the exception of the first 19 TOAs and the last 10 TOAs).

This method, however, causes some events to be detected multiple times. This happens because the effects of $\Delta\nu$ and $\Delta\dot{\nu}$ may be detectable not only in the set of TOAs starting immediately after the event but also in some of the neighbouring trials. Close inspection of the results shows that detections typically cluster in groups of 2–5 trials, separated by no more than 2 days, and that clusters are typically 20 to 30 days apart. To remove the repeated detections and produce a final list of GCs we select from each cluster of candidates the detection with the largest $\Delta\nu$ value. This is a conservative choice which makes the final list of GCs a representation of the maximum possible activity present in the data. Also, this choice follows the experience gained from the detection of previously known glitches (section 5.1).

5 RESULTS

We ran the detector over the 42-ft dataset, using the data from January 1984 to February 2013. The detector found all but one of the known glitches in this time-span as well as a large number of GCs.

5.1 The output of the glitch detector

The only previously known glitch that was not found by the detector occurred on MJD $\sim$ 52146.8, only 63 days after the previous glitch. It was not labelled as a possible glitch because none of the fits, neither the quadratic nor the linear, gave a reduced χ^2 less than 15, one of the conditions to create a GC. The smallest reduced χ^2 among the fits around this glitch was 18. We attribute these poor fits to the influence of the recovery from the previous glitch.

[2] Changing this threshold to 20, for example, we obtained 12 new GCs, homogeneously distributed across the frequency range of GCs. There are no effects on the statistical results described in later sections.

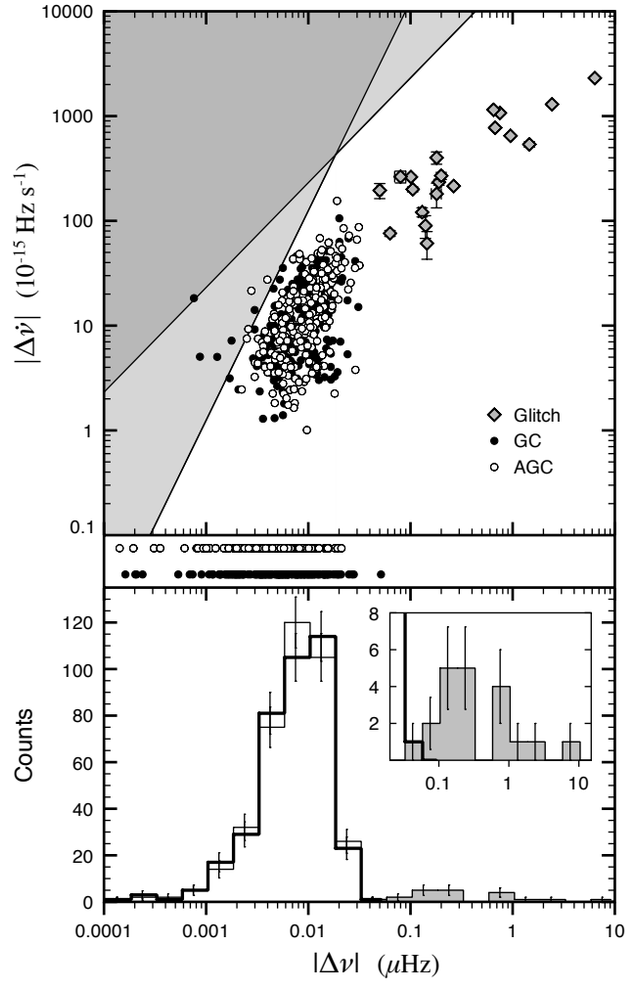

Figure 2. Previously known glitches (diamonds, from Espinoza et al. (2011b,a)), glitch candidates (GCs) and anti-glitch candidates (AGCs). The top panel shows their distribution in the $|\Delta\nu|$–$|\Delta\dot{\nu}|$ plane. The straight line with the smallest slope represents the detection limit expected (Eq. 2) for a cadence of $\Delta T = 1$ day. The one with the largest slope represents the detection limit expected for residuals of $\sigma_\phi = 0.0004$. Our observations are not sensitive to glitches in the shaded areas above these lines. The middle panel shows the $|\Delta\nu|$ values of those candidates with undetectable $|\Delta\dot{\nu}|$. The lower panel shows histograms for the $|\Delta\nu|$ values of the known glitches (filled grey), GCs (thick black) and AGCs (thin black). The inset shows a zoom in the region $|\Delta\nu| > 0.02\,\mu$Hz.

The sizes $\Delta\nu$ and $\Delta\dot{\nu}$ that the detector measured for the known glitches are in good agreement with the values published by Espinoza et al. (2011b). Nevertheless, some differences can be found among the $\Delta\nu$ measurements. As discussed above, because of the way the detector works, the effects of every glitch were detected in more than one set of TOAs. The $\Delta\nu$ value coming from the set of TOAs offering the best fit (smallest reduced χ^2) is always smaller than the published value, which is obtained by standard timing techniques. In addition, this set of TOAs is normally the one starting one to three TOAs after the glitch epoch. On the other hand, the glitch sizes obtained when testing at the correct glitch epoch are typically the largest and the most similar to the published values, though the fits have

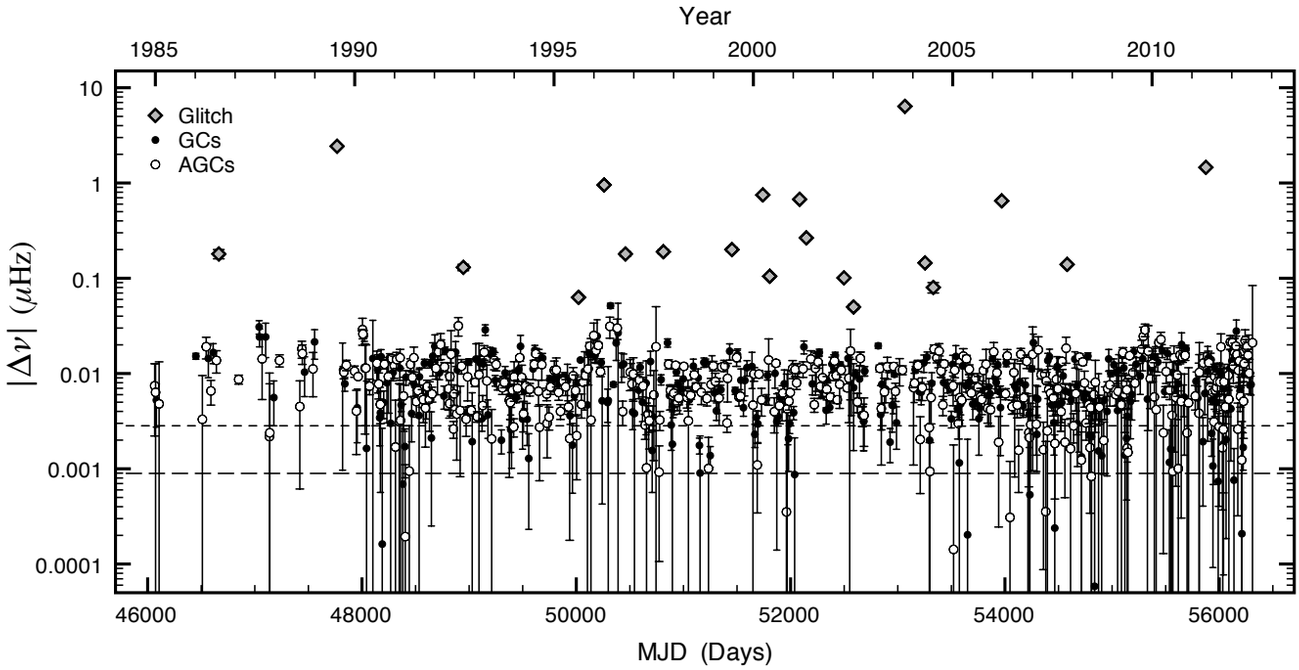

Figure 3. Time sequence of glitches, glitch candidates (GCs) and anti-glitch candidates (AGCs). Horizontal lines indicate the detection limits (Eq. 2) for $|\Delta\dot{\nu}| = 1 \times 10^{-15}\,\mathrm{Hz\,s^{-1}}$ (long dashed) and $10 \times 10^{-15}\,\mathrm{Hz\,s^{-1}}$ (short dashed). The low number of detections between the years 1985 and 1990 is caused by the presence of short gaps with no observations.

larger reduced χ^2 values. These effects are likely caused by unmodelled rapid exponential recoveries and were taken into account when selecting one candidate from a group of several candidates in the overall search. The uncertainties of the GC sizes are the square root of the variances of the parameters, given by the Levenberg-Marquardt algorithm used to fit the data, multiplied by the square root of the reduced χ^2 of the fit.

We reviewed the output of the glitch detector with the aim of producing a clean list of GCs. First, we removed from the original list all those GCs related to known glitches. Then we kept only one candidate (the one having the largest $\Delta\nu$) per event, as mentioned in the description of the method (section 4.1). Next, we visually inspected the timing residuals for all GCs having $\Delta\nu \geq 0.02\,\mu$Hz and eliminated three which involved large data gaps or with timing residuals clearly contaminated by glitch recoveries. We also examined the possibility that some GCs in this $\Delta\nu$ range could be caused by rapid changes in the electron density towards the Crab pulsar (Lyne et al. 1993), which strongly affects the travel time of the pulsar emission at these low frequencies, introducing signatures in the data which can mimic a glitch. To do so, we used observations taken at higher frequencies (mostly at 1400 MHz, with the Lovell telescope) and removed a further three GCs that were clearly caused by this effect. However, the cadence of the Lovell observations is not as rapid as that of the 42-ft observations and we were unable to confirm some other possible cases of such non-achromatic events. We inspected the timing residuals of the largest remaining GCs ($|\Delta\nu| \geq 0.02\,\mu$Hz) and found their signatures to be indistinguishable from timing noise, though we acknowledge that discrimination between small glitches and timing noise is difficult. Nonetheless, in many cases no sharp transitions, typical of the known glitches, are observed at the GC epochs and the residuals are consistent with a smooth connection with the pre-GC-epoch residuals.

Our final list contains 381 GCs. They are homogeneously distributed over the entire time-span and are clustered as a population in $\Delta\nu$–$|\Delta\dot{\nu}|$ space (Figs. 2, 3). The vast majority of them exhibit $\Delta\nu$ steps that are smaller than all previously detected glitches, leaving a gap between the $\Delta\nu$ distributions of real glitches and GCs which would be hard to populate with undetected events.

5.2 Search for anti-glitches

Given the distinct properties of the GC population, it is possible that the glitch-like signatures found by the detector are a component of the Crab pulsar's timing noise. To test this idea and explore the noise nature of these irregularities, we performed a search for events with the opposite signature to a glitch, i.e. anti-glitch candidates (AGCs) with $\Delta\nu < 0$ and $\Delta\dot{\nu} \geq 0$, which are subject to the equivalent detection constraints as the normal glitches. After removing repeated detections and 10 events caused by glitch recoveries, gaps with no data and non-achromatic events (see above), we obtain 383 AGCs. They are also separate from the glitch population and show very similar characteristics to the GCs (Figs. 2, 3).

6 DISCUSSION

6.1 GCs and AGCs: glitches or timing noise?

Using the Kolmogorov–Smirnov (K–S) test we can compare the $|\Delta\nu|$ distributions for GCs and AGCs, which are found to be statistically consistent with coming from the same parental distribution (with a K–S statistic of $D = 0.037$ and $p_{KS}(D) = 0.96$, thus a probability of only $\sim 4\%$ for a false null hypothesis). The $|\Delta\nu|$ distributions for GCs and AGCs can be well described by lognormal distributions, with probability density function (PDF) of the form

$$p(|\Delta\nu|) = \frac{1}{\sqrt{2\pi}\sigma(|\Delta\nu| - \theta)} \exp\left[-\left(\frac{\ln(|\Delta\nu| - \theta) - \mu}{4\sigma}\right)^2\right] \quad (3)$$

for $|\Delta\nu| > \theta$, which gives $p_{KS} = 0.85$ and 0.92 respectively. However, this result is only indicative since the lower ends of these distributions are not well probed by the observations (Figs. 2, 3).

We also compared the $|\Delta\nu|$–$|\Delta\dot{\nu}|$ distributions of GCs and AGCs using a 2-dimensional K–S test (Press et al. 1992). The test gives $D_{2D} = 0.084$, implying a probability of $\sim 40\%$ that they come from the same distribution. This relatively low probability is likely to be produced by differences in the $|\Delta\dot{\nu}|$ distributions between GCs and AGCs, since a K–S test over these two gives $p_{KS} = 0.55$, considerably smaller than the one for $|\Delta\nu|$.

Neither a power-law nor a lognormal distribution can describe well the joint $\Delta\nu$ distribution of the 20 glitches plus all the GCs, with $p_{KS} < 10^{-4}$. A power-law with a lower cut-off at $\Delta\nu \sim 0.01\,\mu$Hz, to account for the incompleteness of the sample at small sizes, gave a similarly poor fit.

Although it is possible that some of the GCs correspond to real glitches, we interpret all the above results as confirmation that the GCs and AGCs are generated by a symmetric noise process and that no new glitches have been found. This timing noise component produces a continuous departure from a simple slow-down trend with variations that can be characterised by changes of $|\Delta\nu| \leq 0.03\,\mu$Hz and $|\Delta\dot{\nu}| \leq 200 \times 10^{-15}\,\mathrm{Hz\,s^{-1}}$.

6.2 The glitch size distribution

Having established that the 20 glitches form the complete sample of glitches the Crab pulsar has had in the last 29 years, we can address their statistical properties.

To determine the best-fit exponent α for a power-law PDF of the form

$$p(\Delta\nu) = C\Delta\nu^{-\alpha} \quad , \quad (4)$$

with $\Delta\nu_{\min} \leq \Delta\nu \leq \Delta\nu_{\max}$ and $C = (1-\alpha)(\Delta\nu_{\max}^{1-\alpha} - \Delta\nu_{\min}^{1-\alpha})^{-1}$, we use the maximum-likelihood estimator method. Setting $\Delta\nu_{\min} = 0.05\,\mu$Hz and $\Delta\nu_{\max} = 6.37\,\mu$Hz, the values for the smallest and largest glitches observed respectively, we obtain $\alpha = 1.36\,(+0.15, -0.14)$ (Fig. 4). The value of the exponent does not depend strongly on the choice of limits, as long as these are a few times smaller or larger than the observed ones. To assess the goodness of the fit, we calculate the K–S statistic, $D = 0.1$, and its probability value $p_{KS}(D) = 0.9$, which corresponds to a 10% probability that our null hypothesis (that the data follow the PDF described by Eq. (4)) is false. Thus our results confirm that the Crab glitch $\Delta\nu$ distribution is consistent with a power-law, a description motivated by theoretical models.

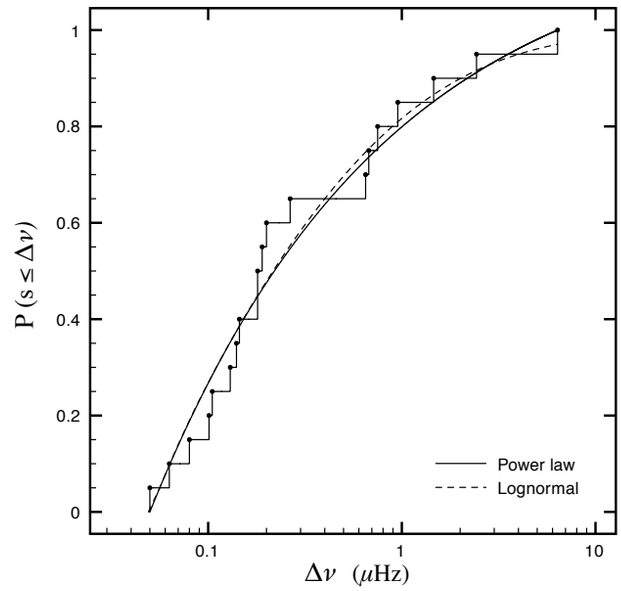

Figure 4. The cumulative distribution function of the observed glitch sizes, s, and the corresponding power-law fit (solid line) given by Eq. 4, with $0.05 \leq \Delta\nu\,(\mu\mathrm{Hz}) \leq 6.37$ and $\alpha = 1.36$. The dashed line corresponds to a lognormal fit (Eq. 3) with $\mu = -1.79$, $\sigma = 1.9$ and $\theta = 0.049\,\mu$Hz.

If this power-law continued below $\Delta\nu_{\min}$, we would expect to have detected more than 10 glitches with $0.02\,\mu\mathrm{Hz} < \Delta\nu < \Delta\nu_{\min}$ in the searched data, and the gap between glitches and GCs (in Figs. 2 and 3) should have been populated. Thus we observe a rapid fall-off of the power-law for $\Delta\nu < \Delta\nu_{\min}$.

However, the small sample size makes it impossible to exclude other distributions. For example, the same K–S probability is obtained for a lognormal distribution (Eq. 3) with parameters $\mu = -1.79$, $\sigma = 1.9$ and $\theta = 0.049\,\mu$Hz, whose probability density function also quickly vanishes for $\Delta\nu < \Delta\nu_{\min}$. The same conclusions hold if the four glitches from before the start of this dataset are included in the sample.

Further confirmation of the rare occurrence of small glitches comes from the study of the $|\Delta\nu|$–$|\Delta\dot{\nu}|$ distribution. Having shown that the latter is not affected by observational biases, the correlation between $|\Delta\nu|$ and $|\Delta\dot{\nu}|$ (apparent in Fig. 2) is confirmed to be a robust feature. While $\Delta\nu$ measurements are very accurate, the acquired values of $\Delta\dot{\nu}$ are less certain and depend upon the method used to determine them, leading sometimes to large discrepancies. For this work we consistently calculated the glitch parameters for all 20 glitches, using the technique described in Espinoza et al. (2011b). Using those measurements, the Spearman's rank correlation coefficient between $\Delta\nu$ and $|\Delta\dot{\nu}|$ is $\mathrm{rs} = 0.776$ with $p(\mathrm{rs}) = 6 \times 10^{-5}$, which indicates a strong correlation. We note that the correlation becomes stronger if the four early glitches are included.

Given this relationship, any additional glitches would occupy a region of the $|\Delta\nu|$–$|\Delta\dot{\nu}|$ space well probed by the observations. Therefore, we observe a rapid decrease of the

probability for glitch sizes below $\Delta\nu_{\rm min}$, which cannot be ascribed to incompleteness of the sample and hence indicates the existence of a minimum glitch size for the Crab pulsar.

7 IMPLICATIONS FOR THEORETICAL MODELS

Such a limit for the smallest glitch size is challenging to our current understanding of glitches and has the potential to constrain the proposed mechanisms.

Some simple considerations can be used to get a rough order of magnitude estimate for the number of neutron superfluid vortices that need to unpin to produce the smallest Crab glitch. Each superfluid vortex carries a quantum of circulation $\kappa = h/2m_n \sim 2\times10^{-3}\,{\rm cm^2\,s^{-1}}$. Neglecting differential rotation of the superfluid (and entrainment), its total circulation at distance r from the rotational axis will be $\Gamma = \oint v_s \cdot dl = N_v(r)\kappa = 2\Omega_s(r)A$, where $N_v(r)$ is the number of vortices in the enclosed area A and Ω_s is the superfluid angular velocity. Using $r = 10^6\,{\rm cm}$, the total number of vortices for the Crab pulsar is of the order $N_v \sim 6\times10^{17}$. Conservation of the total angular momentum implies that the angular velocity change of the superfluid, $\delta\Omega_s$, relates to the observed glitch size $\Delta\nu$ by $\delta\Omega_s = 2\pi\Delta\nu I_c/I_s$, where I_c is the moment of inertia of the coupled component and I_s is the superfluid moment of inertia that participates in the glitch. Using a typical value of $I_c/I_s \sim 10^2$ and $\Delta\nu = \Delta\nu_{\rm min} = 0.05\,\mu{\rm Hz}$, the total number of vortices must be reduced by $\delta N_v \sim 10^{11}$.

The actual change in the superfluid angular momentum L_s depends on the number of vortices that unpinned, the location and size of the region where this happened and the distance travelled by those vortices before they repin. For a more rigorous estimate, the change in L_s can be approximated by $\Delta L_s \sim \hat{\rho}\kappa\xi R^3\delta N_v$, where $\hat{\rho}$ is the average density of the region involved, R is the stellar radius and ξ is the fraction of R that unpinned vortices travel (Warszawski & Melatos 2013). For a typical value of $I_c \sim 10^{45}{\rm g\,cm^2}$ for the moment of inertia of the coupled component, the smallest glitch observed translates to an angular momentum change of $\Delta L_c = 2\pi I_c\Delta\nu_{\rm min} = 3\times10^{38}\,{\rm g\,cm^2\,s^{-1}}$. Conservation of angular momentum leads to $\xi\delta N_v \simeq 1.5\times10^9$ if one assumes typical values for the base of the crust, like $R = 10\,{\rm km}$ and $\rho = 10^{14}\,{\rm g\,cm^{-3}}$. Vortices are expected to repin after encountering a few available pinning sites, however as a conservative order of magnitude estimate we assume they cover a distance comparable to the thickness of the crust (1.5 km) and take $\xi \leq 0.15$, which means that at least 10^{10} vortices must unpin in a glitch with $\Delta\nu = \Delta\nu_{\rm min}$. Therefore the observed minimum glitch size, which is well above that expected for single-vortex unpinning events, implies the existence of a smaller characteristic length-scale which sets the lower cut-off for the range of the scale-invariant behaviour.

The vortex avalanche model is based on the notion of self-organised criticality (SOC, Bak et al. 1987), applications of which can be found for example in earthquake dynamics (Hergarten 2002) or superconducting flux-tube avalanches (Wijngaarden et al. 2006). SOC occurs without the need of fine tuning of parameters, in several dynamical systems consisting of many interacting elements (the superfluid vortices in the case of a neutron star) which, under the act of an external slow driving force (the spin-down of the star), self-organise in a critical stationary state with no characteristic spatiotemporal scale. A small perturbation in such systems can trigger an avalanche of any size. Thus in the glitch avalanche model of Warszawski & Melatos (2008) vortex density is assumed to be greatly inhomogeneous and many metastable reservoirs of pinned vortices are formed, which relax independently giving rise to the observed spin-ups. Since such a system has no preferred scale the resulting glitch magnitudes follow a power-law distribution. This behaviour should however continue down to events involving the unpinning of only a few vortices, which is orders of magnitude below the observed cut-off.

The coherent noise model (Newman & Sneppen 1996) is a different, non critical mechanism which produces scale-free dynamics, even in the absence of interaction between the system's elements. In such systems a global stress is imposed to all elements coherently, to which they respond if it exceeds their individual unpinning threshold, giving rise to avalanches of various sizes. Both threshold levels (for each element) and stress strength are randomly chosen from respective probability distribution functions. The elements with thresholds smaller than the applied stress will participate in an avalanche and then be re-assigned new threshold values. New thresholds must always be assigned to a few elements, even when no avalanche is triggered, otherwise such a system will stagnate. A possible mechanism for this process in superfluids is the thermally activated unpinning of vortices (Melatos & Warszawski 2009), while the global Magnus force acts as the coherent stress. The model predicts a minimum for the glitch magnitude, which represents the thermal creep only events, present even if all thresholds lie above the applied stress strength. But it also predicts an excess (with respect to the resulting power-law) of such small glitches, in contradiction to what we observe for the Crab pulsar. The lack of this overabundance of small events requires a broad distribution for the pinning potentials. Melatos & Warszawski (2009) studied the top-hat distribution and applied their model to the Crab pulsar. They found that the half-width of the distribution should be comparable to the mean pinning strength. Even when such a broad distribution of pinning energies is introduced, independent unpinning of vortices as a random Poisson process of variable rate proves insufficient to produce scale-invariant glitches (Warszawski & Melatos 2013), indicating that the interaction between vortices and collective unpinning (a domino-like process) must be taken into account. The most prominent mechanism for collective unpinning is the proximity effect, in which a moving vortex triggers the unpinning of its neighbours. However such a mechanism requires extreme fine tuning, since power-law size distributions occur only if this effect is neither too weak (where thermal creep dominates) nor too strong (which always leads to large, system-spanning, avalanches) (Warszawski & Melatos 2013).

Another process which could lead to scale-invariant glitches are crustquakes (Morley & Schmidt 1996). Stresses develop in the solid crust of a neutron star because of the change in its equilibrium oblateness as the spin decreases, but also due to the interaction of the crustal lattice with the magnetic field and superfluid vortices in the interior. If the crust cannot readjust plastically it will do so abruptly when the breaking strain $\epsilon_{\rm cr} = \sigma_{\rm cr}/\mu$ is exceeded (where

$\sigma_{\rm cr}$ is the critical stress and μ the mean modulus). This will result in both a spin-up (due to the moment of inertia decrease) and in a reaction of the superfluid (Alpar et al. 1996; Ruderman et al. 1998), which is evident in the post-glitch relaxation. The maximum fractional moment of inertia change associated with the $\Delta\nu_{\rm min}$ glitch is $|\Delta I|/I \leq 10^{-9}$; we note here that the glitch size can be significantly boosted by the crustquake induced unpinning of vortices (Larson & Link 1999; Eichler & Shaisultanov 2010). Elastic stress on the crust due to change of the equilibrium oblateness builds up because of the almost-constant secular $\dot{\nu}$. Therefore the critical stress $\sigma_{\rm cr}$ will be reached in regular time intervals if all stress is relieved in each crustquake, and the total energy released will be $\Delta E_{\rm el} \propto \epsilon_{\rm cr}^2$. If the stress is only partially relaxed then the energy released will depend on the stress drop $\Delta\sigma$, and the time interval to the next crustquake will depend on the size of the preceding one. The latter correlation is observed for the glitches in PSR J0537-6910, which have been interpreted as crustquakes (Middleditch et al. 2006). For the Crab pulsar however, the lack of any such trends in our glitch sample indicates a more complicated picture.

8 CONCLUSIONS

We have quantified our current glitch detection capabilities and, after a meticulous search for small glitches, we have shown that in the case of the Crab pulsar all glitches in this dataset have already been detected. The full glitch size distribution exhibits an under-abundance of small glitches and implies a lower cut-off at $\Delta\nu \sim 0.05\,\mu$Hz. The existence of such a minimum glitch size implies a threshold-dominated process as their trigger, which still needs to be identified.

Besides the occasional glitches, we have detected a continuous presence of timing noise having a well defined maximum amplitude, which can be described by step changes $|\Delta\nu| \leq 0.03\,\mu$Hz and $|\Delta\dot{\nu}| \leq 200 \times 10^{-15}\,{\rm Hz\,s^{-1}}$. The distinct properties of this noise component compared to the glitches imply that timing noise cannot be attributed solely to unresolved small glitches produced by the exact same mechanism.

ACKNOWLEDGMENTS

Pulsar research at JBCA is supported by a Consolidated Grant from the UK Science and Technology Facilities Council (STFC). C.M.E. acknowledges the support from STFC and FONDECYT (postdoctorado 3130512). D.A. and A.L.W. acknowledge support from an NWO Vidi Grant (PI Watts).

REFERENCES

Alpar M. A., Chau H. F., Cheng K. S., Pines D., 1993, ApJ, 409, 345
Alpar M. A., Chau H. F., Cheng K. S., Pines D., 1994, ApJ, 427, L29
Alpar M. A., Chau H. F., Cheng K. S., Pines D., 1996, ApJ, 459, 706
Anderson P. W., Itoh N., 1975, Nature, 256, 25
Andersson N., Glampedakis K., Ho W. C. G., Espinoza C. M., 2012, Phys. Rev. Lett., 109, 241103
Andersson N., Glampedakis K., Hogg M., 2013, Phys. Rev. D, 87, 063007
Bak P., Tang C., Wiesenfeld K., 1987, Physical Review Letters, 59, 381
Baym G., Pethick C., Pines D., 1969, Nature, 224, 673
Baym G., Pines D., 1971, Ann. Phys. (U.S.A.), 66, 816
Chamel N., 2013, Phys. Rev. Lett., 110, 011101
Eichler D., Shaisultanov R., 2010, ApJ, 715, L142
Espinoza C. M., Jordan C., Bassa C., Janssen G., Lyne A. G., Smith F. G., Stappers B. W., Weltevrede P., 2011a, The Astronomer's Telegram, 3777, 1
Espinoza C. M., Lyne A. G., Stappers B. W., Kramer M., 2011b, MNRAS, 414, 1679
Glampedakis K., Andersson N., 2009, Phys. Rev. Lett., 102, 141101
Gullahorn G. E., Isaacman R., Rankin J. M., Payne R., 1977, AJ, 82, 309
Haensel P., Potekhin A. Y., Yakovlev D. G., 2007, Astrophysics and Space Science Library, Vol. 326, Neutron Stars 1 : Equation of State and Structure, Haensel P., Potekhin A. Y., Yakovlev D. G., eds. Springer
Haskell B., Pizzochero P. M., Sidery T., 2012, MNRAS, 420, 658
Hergarten S., 2002, Self-Organized Criticality in Earth Systems. Springer, Berlin
Hobbs G. B., Edwards R. T., Manchester R. N., 2006, MNRAS, 369, 655
Larson M. B., Link B., 1999, ApJ, 521, 271
Link B., Epstein R. I., 1996, ApJ, 457, 844
Lyne A., Hobbs G., Kramer M., Stairs I., Stappers B., 2010, Science, 329, 408
Lyne A. G., Pritchard R. S., Smith F. G., 1988, MNRAS, 233, 667
Lyne A. G., Pritchard R. S., Smith F. G., 1993, MNRAS, 265, 1003
Melatos A., Peralta C., Wyithe J. S. B., 2008, ApJ, 672, 1103
Melatos A., Warszawski L., 2009, ApJ, 700, 1524
Middleditch J., Marshall F. E., Wang Q. D., Gotthelf E. V., Zhang W., 2006, ApJ, 652, 1531
Morley P. D., Schmidt I., 1996, Europhysics Letters, 33, 105
Newman M. E. J., Sneppen K., 1996, Physical Review E, 54, 6226
Pizzochero P. M., 2011, ApJ, 743, L20
Press W. H., Teukolsky S. A., Vetterling W. T., Flannery B. P., 1992, Numerical Recipes: The Art of Scientific Computing, 2^{nd} edition. Cambridge University Press, Cambridge
Radhakrishnan V., Manchester R. N., 1969, Nature, 222, 228
Ruderman M., 1969, Nature, 223, 597
Ruderman M., Zhu T., Chen K., 1998, ApJ, 492, 267
Wang J., Wang N., Tong H., Yuan J., 2012, Astrophys. Space Sci., 340, 307
Warszawski L., Melatos A., 2008, MNRAS, 390, 175
Warszawski L., Melatos A., 2013, MNRAS, 428, 1911
Wijngaarden R. J., Welling M. S., Aegerter C. M., Menghini M., 2006, European Physical Journal B, 50, 117
Wong T., Backer D. C., Lyne A., 2001, ApJ, 548, 447